

UNITED STATES PATENT APPLICATION FOR


**DEPOPULATED
PROGRAMMABLE LOGIC ARRAY**

Inventors:
Robert Osann, Jr.
Patrick Hallinan
Jung Lee
Shridhar Mukund

**CERTIFICATE OF MAILING BY "EXPRESS MAIL"
UNDER 37 C.F.R. § 1.10**

"Express Mail" mailing label number: EL670727865US
Date of Mailing: April 5, 2001

I hereby certify that this correspondence is being deposited with the United States Postal Service, utilizing the "Express Mail Post Office to Addressee" service addressed to **Box PATENT APPLICATION, Assistant Commissioner for Patents, Washington, D.C. 20231** and mailed on the above Date of Mailing with the above "Express Mail" mailing label number.



Matthew A. Mahling
Signature Date: April 5, 2001

**DEPOPULATED
PROGRAMMABLE LOGIC ARRAY**

Inventors:
Robert Osann, Jr.
Patrick Hallinan
Jung Lee
Shridhar Mukund

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims priority to Provisional Application No. 60/231,059, filed September 8, 2000. This application is related to Application Serial No. 09/512,783, filed February 25, 2000, entitled "Programmable Logic Array Embedded in Mask-Programmed ASIC," incorporated by reference herein.

FIELD OF INVENTION

[0002] This invention generally relates to programmable logic arrays (PLAs), and particularly to minimizing the size and speed of PLAs.

BACKGROUND

[0003] Programmable Logic Arrays (PLAs) have been used in the semiconductor industry for many years. They are essentially generic electronic circuits, able to be programmed to provide a variety of logic functions. PLAs are well understood in the art, but a brief explanation is provided below.

[0004] In PLAs, a logic function to be implemented is expressed in product terms (AND terms) and sum terms (OR terms). Each product term is generated by circuitry that can be programmed to form the AND of any subset of the inputs to the PLA and their complements. Product terms are then be summed, or OR'd, through additional programmable circuitry.

[0005] More specifically, PLAs are generally constructed in the form of regular arrays, with the input lines being orthogonal to the product lines, as shown in the generalized circuit diagram of Fig. 1. Fig. 1 shows one type of PLA known as a PAL. In Fig. 1, PAL 100 includes inputs 102, A-D, where each input and its complement is input into programmable AND array 104 on lines 106. AND lines 108 are formed orthogonal to lines 106. When a selected intersection between an input-term line 106 and an AND line 108 is programmed, then "AND terms" or "product terms" 110 are formed. The product terms 110 are then fed into OR gates 112, forming "sum terms" 114.

[0006] As is understood in the art, product-terms are often implemented with a wired-OR mechanism, where at each intersection of lines 106 and 108 programmable transistors 115 are used. A pull-up 116 is also used, as shown in the generalized circuit diagram of Fig. 2. Although the pull-up 116 is shown as a resistor in Fig. 2, frequently a passive pull-up is implemented with a biased P-channel transistor instead. It will be noted that when using the structure of Fig. 2, the AND gates 109 shown in Fig. 1 are not actually physically present, but instead in Fig. 1 AND gates 109 represent logically the function being performed.

[0007] A second type of PLA structure (sometimes referred to as a "Full PLA") is shown in Fig. 3, having both a programmable AND array 104 and a programmable OR array 117. In other words, both product terms and sum terms can be programmed. Thus, the device of Fig. 3 can implement any set of combinational logic limited only by the number of inputs, outputs, and product terms.

[0008] Relative to the device of Fig. 1, the device of Fig. 3 sacrifices some speed, but has greater programming flexibility and is better for implementing state machines. As well, unlike the Fig. 1 device, in the Fig. 3 device, each product term can be used by, or shared among, all sum terms. Still, as shown in Fig. 4, some PLA devices create OR arrays where selected groups of the product terms are shared among selected groups of the sum terms.

[0009] Flexibility of PLA structures, including those of Figs. 1 and 3, can be further enhanced by adding flip-flops 118 (shown in Fig. 3) to one or more of the

outputs to create general-purpose sequential circuits. As shown in Fig. 5, inverted product terms such as !PT can also be fed back directly into the AND array. As used herein, complemented terms will be indicated with either a "!" or with an overbar, as is the practice in the art. As shown in Fig. 5a, when many product terms are inverted and fed back, it forms what is known as a NAND-NAND (or NOR-NOR) array.

[0010] Many PLA devices also include post-array logic structures as shown in Fig. 6. The most frequent need for such structures is to complement a logic function – otherwise, to complement the function, one would have to resort to DeMorgan's Law ($!(AB) = (!A + !B)$), typically resulting in a large number of product terms that would need to be programmed. XOR gates 120 are commonly used as such post-array structures. Inversion can be controlled by (1) a sum term representing a complex function (as for XOR gate 120₁), (2) a single product term (as for XOR gate 120₂), or (3) by a single logic value that can be programmed to a logical "1" (V_{dd}) or "0" (V_{ss}) (as for XOR gate 120₃).

[0011] Similarly, pre-array logic structures 122 are also often used as shown in Fig. 7. Such pre-array logic structures are often useful in performing arithmetic functions such as accumulators. Both post-array and pre-array logic structures are used to reduce the number of product terms that would otherwise be required, and hence allow the use of smaller and faster PLA devices.

[0012] Each of the conventional PLA structures described in Figs. 1-7 is usually a discrete device and is programmable and reprogrammable by the user either using a specialized programming device or in-system as is understood in the art. As should be understood in the art from Figs. 1-7, there are numerous types of PLAs available and those described are exemplary only.

[0013] As alluded to above with the discussion of pre- and post-array structures, it is generally desirable to use the smallest PLA device available, for real-estate reasons as well as speed. To that end, "PLA folding" has been studied and is described in Ferreira and Song, "Achieving Optimality for Gate Matrix Layout and PLA Folding: a Graph Theoretic Approach", 1992. In general, PLA folding is a

technique that, given a PLA, attempts to produce an equivalent PLA that occupies less space. For example, consider the following logic functions to be implemented by a PLA:

5

$$f_6 = \bar{v}_1 v_5 + \bar{v}_2 \bar{v}_3 + v_1 v_3 + \bar{v}_5$$
$$f_7 = v_2 v_4 + \bar{v}_2 \bar{v}_4$$

These functions are shown implemented in a traditional PLA in Fig. 8, where the “dots” represent a programmed connection. Through a series of matrix manipulation and rearranging of rows, the result of the folding is shown in Fig. 9. More details of this example can be found in Ferreira and Song (cited above). As shown in Fig. 9, the folded PLA is much more compact, retaining as much physical regularity as possible by causing the tiling of product term groupings.

[0014] Still, this PLA folding technique has gained little practical application. To effectively use folding, an engineer must know with relative certainty the functions to be programmed. And the freedom to make changes to the functionality is severely limited by the “folds.” Moreover, PLAs are typically only available as discrete devices (e.g., PLDs) in fixed sizes with fixed characteristics. So while folding techniques might be useful in theory, there are essentially no devices using them.

SUMMARY

SUB 7
B1

[0015] A PLA in accordance with the invention is designed after knowing the initial logic to be implemented by the PLA. Once such logic is known, a PLA sized specifically to that logic is modelled. Then unnecessary programmable connections are removed creating a relatively small, but fast PLA to implement the given function. Additional programmable connections are then re-added to the model to allow for future reprogramming. Finally, the modelled PLA is constructed.

[0016] The result is a PLA that achieves a maximum amount of depopulation while still implementing a logic function and maintaining flexibility for future

reprogramming. In addition a PLA in accordance with the invention can be built so that no matter what functionality is programmed, the performance characteristics for the device remain the same.

5 [0017] A PLA in accordance with the invention does not require a regular array structure. Therefore, such a PLA is flexible for place and route considerations, particularly when integrated with other logic in larger circuits.

BRIEF DESCRIPTION OF THE DRAWINGS

10 [0018] The present invention is described with respect to particular exemplary embodiments thereof and reference is accordingly made to the drawings in which:

[0019] Fig. 1 is a functional diagram of a conventional PAL-type PLA;

[0020] Fig. 2 illustrates a conventional wired-OR structure;

[0021] Fig. 3 is a functional diagram of a conventional full PLA;

[0022] Fig. 4 illustrates product-term sharing in conventional PLA structures;

15 [0023] Fig. 5 illustrates complemented product terms in conventional PLA structures;

[0024] Fig. 5a illustrates a conventional NAND-NAND array;

[0025] Fig. 6 illustrates post-array logic in conventional PLA structures;

[0026] Fig. 7 illustrates pre-array logic in conventional PLA structures;

20 [0027] Figs. 8-9 illustrate an example of PLA folding;

[0028] Fig. 10 is a functional diagram of a fully populated PLA in accordance with an embodiment of the invention;

[0029] Fig. 11 is a functional diagram of an intermediately populated PLA in accordance with an embodiment of the invention;

25 [0030] Fig. 12 is a functional diagram of a minimally populated PLA in accordance with an embodiment of the invention;

[0031] Figs. 13 and 13a functionally illustrate a core cell in accordance with an embodiment of the invention;

30 [0032] Figs. 14 and 14a functionally illustrate a core cell in accordance with another embodiment of the invention;

[0033] Fig. 15 is a functional diagram of an embodiment of the invention;

[0034] Fig. 16 illustrates an embodiment of the invention with pre- and post-array logic;

5 [0035] Figs. 17-21 illustrate complemented P-terms in accordance with various embodiments of the invention.

DETAILED DESCRIPTION

10 [0036] The beauty of a PLA structure is that any logic function can be implemented that fits within the size of the structure, where the size is determined by the number of product terms (p-terms), sum terms, inputs, registered outputs, and feedback terms. The flexibility and programmability of a PLA is especially desirable for circuitry that is not fully verified or may need to be changed in the future. However, the genericness of the PLA can cause the size of the PLA to be significantly large since the physical paths and gates to support all possible overall
15 functionalities must be constructed. In fact, a programmed PLA usually has only a small percentage of the possible logic functionality utilized. In addition to the large size of a conventional PLA relative to the logic it is implementing, the circuit speed is also hindered due to the additional capacitive loading of driving circuitry, as well as the capacitance of driving longer lines due to the larger overall size of
20 the array.

[0037] But if the function of the PLA is known before the PLA is physically constructed then a methodology can be employed where the portions of the PLA that are not required to perform the specific function can be eliminated. Even with this “depopulation” in effect, changes and enhancements can still be implemented
25 if some amount of additional product term capability is included in the array. An example of a PLA built after its initial functionality is known can be found in Application Serial No.09/512,793, filed February 25, 2000, and entitled “Programmable Logic Array Embedded in Mask-Programmed ASIC,” incorporated by reference herein, which describes a reprogrammable PLA implemented in an
30 otherwise non-reprogrammable ASIC.

SUB
B2

[0038] Thus, in accordance with the invention, selected programmable connections (e.g., intersections) of a PLA are removed, “depopulating” the PLA, to minimize the space and maximize the speed of the PLA. The connections that are removed are connections that will not be used in by the functionality to be initially programmed in the PLA. Nonetheless, certain initially unnecessary programmable connections are selectively placed in the array to allow for future reprogramming of the array with functionality different then that initially intended.

Depopulation

[0039] Fig. 10 shows a “fully populated PLA” 200. Each intersection in the AND array 204 between input lines 206 and AND lines 208 includes a programmable connection, here shown as core cells 203, which will be described in further detail below and are marked with a “Q”. Similarly, each intersection in the OR array 217 has a similar core cell 203. Buffers 220 may be used in some embodiments for driving the various signals. But because of the number of potential programmable signals, a single buffer in each column may not be sufficient and, therefore, buffer trees may be required to sufficiently drive the loads.

[0040] The PLA 200 is designed and constructed specifically to accommodate a function defined by a user. To do so, the user provides a description of the functionality in HDL or other similar language. The description is analyzed, in some embodiments by software (e.g., a PLA compiler) , and a PLA sized to accomodate the functionality results. In other words, a fully programmable PLA having the minimal number of inputs, outputs (sum terms), and p-terms to accomodate the logic is designed. A small percentage of additional capacity (a small amount of additional p-terms or sum-terms) to handle reprogramming changes are then added, resulting in the design of the PLA structure of Fig.10.

[0041] Fig. 11 shows the PLA 200 redesigned to have a depopulated OR array 217, sometimes referred to herein as an “intermediately populated PLA.” Many of the programmable connections (core cells 203) have been removed (depopulated) as denoted by an “X” at the intersection, indicating that there is no connection —

programmable or otherwise — at that location. The intersections to be depopulated are selected based on the description of functionality to be initially implemented by a PLA as provided by the user. Those intersections that are not required by the user's defined functionality are removed. In addition, where product terms (p-terms) are not shared by OR gates, such as PT1 and PT2, the programmable connection has been removed and a fixed, or hardwired, connection 222 has been inserted instead. Note that an OR array that has been fully depopulated, where all the programmable connections have been removed and replaced with fixed connections, results in a fixed OR array structure such as that found in traditional PAL devices.

[0042] When the programmable connections are removed the PLA size is decreased: not only are the devices for the programmable connections removed, but the OR trees 212 (discussed further below) are also reduced in size. As a result, the overall speed of the device increases.

[0043] But some programmable connections have been strategically kept in Fig. 11. For instance, where a product term is shared between two or more OR gates, such as PT3, then programmable connections remain at the used intersections: if a p-term that was once shared later changes so that it is no longer part of a sum term, it can be programmably disconnected from that sum term. In addition, programmable connections remain in the OR array for spare p-terms such as PT4 to accommodate future programming. Spare p-terms will not generally be used by the user's initial functionality but are provided to add later programming flexibility.

[0044] Fig. 12 shows the PLA 200 further redesigned to have both a depopulated OR array 217 and a depopulated AND array 204, sometimes referred to herein as a "minimally populated PLA." Compared to the structures of Figs. 10 and 11, the configuration of Fig. 12 produces the smallest and fastest PLA, but will sacrifice some of the flexibility found in the other two structures. Depopulation is again determined based on the functionality initially to be implemented by the PLA as defined by a PLA user – programmable connections are removed except where input variables are actually used in the initial logic function to be implemented.

Then selected intersections are re-populated before the array is actually constructed according to a combination of strategies.

5 [0045] One strategy for selective re-population of the AND array is to maintain full population on all columns 224, 226 where input variables are fed back from flip-flops 218 or other registers (e.g., state bits or other registered signal used to construct sequential logic functions). For instance, in sequential logic, many times a future change will involve altering an output to change the edge of a waveform so that the edge happens on a different state. Such alterations can often be easily accomplished by choosing a different logic condition on the feedback bits, but to 10 so with an existing p-term, the needed feedback signals must be populated.

[0046] Another strategy is to allow the PLA user to specify that certain input variables (certain columns) must be fully populated. In other embodiments, the user might specify that a certain input variable be fully populated for certain p-terms or all p-terms that are connected with specific outputs (sum terms).

15 [0047] In addition, any spare p-terms such as PT4 are fully-populated. Since future functionality to be programmed is unknown, it is best to allow full programmability of spare p-terms.

[0048] Strategies for selective re-population of the OR array will be similar to those discussed for the intermediately populated PLA. Other strategies for selective re-population of either array may also be used in other embodiments and those 20 discussed are exemplary only .

SUB B37
25 [0049] In sum, in accordance with the invention, a description of functionality to be implemented in a PLA is received by a program such as a PLA compiler. The description is analyzed and a fully populated PLA is modeled. Then, either the OR array, the AND array, or both are depopulated to remove all unnecessary programmable connections in the model. Strategically selected intersections are then repopulated within the model with programmable connections to allow for future reprogramming. Finally, the PLA is physically constructed based on the modelled PLA.

30

Core Cells

5 [0050] As mentioned above, core cells 203 are programmable connections. The core cells, in one embodiment of a PLA in accordance with the invention, may look much like conventional PLAs as described with respect to Fig. 2 – e.g., product terms implemented with a wired-OR mechanism by connecting multiple transistors to an AND line. In other words, each core cell is essentially a programmable transistor. This wired-OR approach, however, typically requires a PLA structure with a regular array, i.e., one that has wired-OR lines that have the same length and loading, because sense amps are usually coupled to the AND line to amplify signals. If the wired-OR lines are not all the same length with the same loading, then the sense amps would necessarily also have to be different, creating complicated design issues.

10 [0051] So in other embodiments of the invention, the p-terms are implemented differently from conventional wired-OR structures. For instance, Fig. 13 shows a portion of an AND array where core cells 203 program the product terms. As shown in Fig. 13a, two latches or two flip-flops can be used to form such a core cell 203. For discussion purposes only, these devices in the core cells will simply be referred to herein as latches, although it is to be understood that they could be flip-flops in various embodiments. The core cell 203 includes latch 230 and latch 232. Each latch 230, 232 is coupled to mux 236. The select line 206 to mux 236 is an input to the AND array. The mux output 238 is input into multi-input AND gate 214 (Fig. 13), whose output forms a product term. AND gate 214 may be formed using a tree of AND gates as will be understood in the art. Although some illustrations of the invention show a single AND line running to an AND gates, it is to be understood that embodiments such as shown in Fig. 13 (and Fig. 14) will actually have multiple lines coupled to AND gates 214.

25 [0052] Each latch 230, 232 is programmed to store a value that dictates how the core cell 203 will behave upon receiving an input signal on a line 206. Programming is done in accordance with Table 1:

[0053]

TABLE 1

Latch 230	Latch 232	Output 238
1	0	Positive (I)
0	1	Negative (\bar{I})
1	1	Don't care
0	0	Disable P-term

5

10

15

[0054] Thus, when latch 230 stores a logical 1 and latch 232 stores a logical 0, the signal input on line 206 will be output on line 238. When latch 230 stores a logical 0 and latch 232 stores a logical 1, the signal input on line 206 will be inverted on line 238. When both latches 230 and 232 store a logical 1, it indicates a “don’t care” condition – no matter what is input on line 206, the product-term will not be affected. But if latches 230 and 232 both store a logical 0, the entire p-term will be disabled (a 0 input into an AND gate always results in a 0 output).

[0055] When used to program the OR array, the core cell is programmed in accordance with Table 2:

[0056]

TABLE 2

Latch 230	Latch 232	Output 238
1	0	Positive (p-term true)
0	1	Negative (p-term complement)
1	1	Force output high
0	0	Disable P-term

20

25

[0057] An alternative embodiment of a PLA core cell 203 is shown in Figs. 14 and 14a. In Fig. 14a, a single latch or flip-flop 240 is coupled to an OR gate 242. The second input 206 (or 206') to the 2-input OR gate 242 is formed from the input signal or its complement. When latch 240 stores a logical 0, the output of OR gate

242 carries the value of the signal input on line 206 (206'). When the latch 240 stores a logical 1, the input is a "don't care" – it will not effect the p-term.

[0058] A comparison of Figs. 13 and 14 illustrates that in Fig. 14, while only one storage device 240 is required, two physical input lines 206, 206' (one positive and one negative) must be run through the array for each signal input into the array, while in Fig. 13 two storage devices are utilized, 230 and 232, but only one physical input line 206 need be run through the array for each input. In addition, in Fig. 13, AND gate 214 is only half as wide as that in Fig. 14. Although both embodiments are useful, Fig. 13 has the advantage in deep-submicron semiconductor technology that the total wire length has been cut in half. Although the fanout-dependent loading in Fig. 13 will be double that of Fig. 14, it is not fanout loading but absolute wire length that is the predominant factor in determining delays in today's semiconductor devices.

[0059] Nonetheless, some embodiments of the invention will use both of the Fig. 13a and Fig. 14a embodiments. For instance, one embodiment will use the cell structure of Fig. 13a in the AND array while using the cell structure of Fig. 14a in the OR array. Because the OR array does not always require all four of the permutations described in Table 2, but only a care/don't care function, using the embodiment of Fig. 14a allows for a smaller OR array.

[0060] Both the embodiments illustrated in Figs. 13 and 14 are advantageous over wired-OR structures such as those of Fig. 2. First, the structures of Figs. 13 and 14 do not require a regular array structure. Therefore, a PLA in accordance with the invention can be made as small as possible without the limitations of a regular array structure. In some situations, portions of the PLA can even be intermixed with other circuitry in a larger custom chip, such as an ASIC described in Application Serial No. 09/512,783, filed February 25, 2000, and entitled "Programmable Logic Array Embedded in Mask-Programmed ASIC." Still, even if all portions of the PLA are located in proximity to one another, these portions can be scrambled. Thus, although a PLA in accordance with the invention appears to have a regular structure in the figures, such regularity is shown only for clarity

SUB 7
B4
of description. In its physical implementation, a PLA in accordance with the invention need not have any regular structure nor resemble its conventional counterparts at all.

[0061] Further, wired-OR structures are notorious for static power consumption since anytime the AND line is pulled low, current flows. But since the structures of Figs. 13 and 14 use gates, static power consumption can be minimized.

Timing and Performance

[0062] Shown in Fig.15 is an example of a minimally populated PLA 300, having both AND and OR arrays depopulated (similar to Fig. 12). The depopulated OR array results in a number of OR trees such as that shown by OR gates 306-312 and 313-315 as will be understood by those of skill in the art. This example PLA 300 has two outputs 302 and 304, where 302 has a large number of p-terms. The longest critical path will usually be in the largest OR-tree, which path is fixed in the intermediately and minimally populated structures. In Fig.15, the longest path must pass through three OR gates, e.g., 306, 310, and 312.

[0063] The shared and spare p-terms go through the core cells 316-319 as illustrated and therefore have a degree of programmability that allows the steering of p-terms to the different outputs (or disabling the p-terms) when making future changes. Although having these core cells in the paths of the spare and shared p-terms adds some delay to the respective path, the overall effect of this delay in the PLA is not significant: the paths for the shared and spare p-terms join the larger OR-tree close to the output stage of the tree. In other words, the additional delay of the core cells occurs in a shorter path when compared with the longer path occurring in the larger OR-tree. In the example of Fig.15, the largest path must pass through three OR gates, but the shared/spare terms only pass through two OR gates, 311, 312, rendering the delay through the programmed core cell 316 or 317 insignificant to the output 302.

[0064] Thus, a timing model for the PLA structure can be constructed when the PLA is built such that all paths (including the shared and spare paths) are analyzed

in determining the maximum and minimum delay paths. So by using such a timing model, all changes that can possibly be made at a later date will result in delays that are within the original specifications for the maximum and minimum delays. Accordingly, using such a method, when the PLA function changes, the timing will not change. This is particularly significant since in many prior art FPGAs the timing is unpredictable when the function of the device changes.

Use of Pre- and Post-array Structures

[0065] Use of pre- and post-array logic 330, 332 is also easily accommodated in an embodiment of the invention, as illustrated in Fig. 16 (showing an intermediately populated structure, where the core cells in the AND array are not illustrated). In one embodiment, when the HDL description of the PLA is received, then the description is analyzed to determine if any fixed logic functions would be better implemented as pre- or post-array logic in order to reduce the size of the PLA. Once determined, the user may choose which functions to implement as such pre- or post-array logic and the PLA is then generated based on those choices. Alternatively, an embodiment of the invention could automatically select the pre- and post-array structures to meet certain design goals, e.g., the smallest structure possible.

Complemented P-terms

[0066] In accordance with an embodiment of the invention, p-terms can also be easily complemented. For instance, as shown in Fig. 17, a product term A can be fed back in the AND array where it connects to all possible p-terms (column 340). The complement of the fed-back p-term can be chosen by the receiving core cell in column 340 in accordance with Table 1. While this provides a flexible solution, driving all possible p-terms may add delay for a signal path that is already long due to making two trips through the array. But since the user logic to be implemented is already known, in an alternative embodiment the column 340 for the fed-back

term can be depopulated as shown in Fig. 18. Accordingly, a NAND-NAND (or NOR-NOR) array can be implemented, shown in Fig. 19.

[0067] In both Figs. 17 and 18, the feedback loop for signal A would be used to complement a p-term that is incorporated into another p-term. But if an entire p-term needed to be complemented, it could simply be passed through an inverter and then be passed on to the OR gate. Of course a core cell 342 can be used in place of an inverter, as shown in Figs. 17 and 18, since it can both invert and disable a p-term programmably. When the entire p-term A is complemented by passing through a core cell 342 and directly into the OR tree 344 (shown as a 4-input OR gate), it should join the OR tree toward its end in order to minimize any added delay, similar to that described in Fig. 15.

[0068] When a portion of a p-term needs to be complemented, it is desirable to complement the p-term without incurring additional delay, e.g., by two full passes through the array if feed back is used, and at the same time to keep the array small by not adding many p-terms to perform the DeMorgan equivalent of the p-term. Thus, in one embodiment of the invention shown in Fig. 20, the product term A representing the portion of a larger p-term to be complemented passes through a core cell 350, which programmably performs the inversion, and then joins the AND-tree 352 of another p-term. Again, by joining the tree 352 towards its end, almost all of the delay due to the core cell occurs in parallel and there is minimal additional delay added, if any.

[0069] Fig. 21 shows still another embodiment that is similar to that of Fig. 20 except that p-term A can be programmably complemented in its entirety by additional core cell 354 and sent to an OR tree 356 and/or it can be complemented and used in another p-term, joining AND tree 352 as in Fig. 17. As illustrated, by using the core cells, a large amount of flexibility is obtained in that the signal can be steered to different outputs, passed in its true state, complemented, or even disabled.

AND Trees/OR Trees

SUB 7
BS 1

[0070] AND trees and OR trees (as well as buffer trees) have been mentioned above. When building a PLA in accordance with the invention these trees can be built on the fly as needed by the PLA compiler (the program that analyzes the input PLA description and then generates a structural netlist for the PLA). Alternatively, AND and OR trees can be limited to a predefined set of gate primitives used for building trees. Such primitives may be designed such that certain pins/inputs for the primitives are favored due to their speed advantage (e.g., those pins/inputs towards the end of the tree), which will be useful in determining where to add in shared/spare p-terms. Still another method for generating trees is to create a library of all possible trees that will be needed using a synthesis tool like Synopsys Design Compiler as is known in the art. The trees would be synthesized for maximum speed and the tree of the required size would be chosen at the time the PLA is constructed. Again, such trees could be designed such that certain pins/inputs are favored due to their speed advantage.

Specialized Embodiments

[0071] A PLA in accordance with an embodiment of the invention could be designed to perform two diverse tasks. In other words, a single PLA could be built, but it may be used in circuitry that could be applied to two distinct applications. In such a case, the functionality for the PLA for each task will be different and the PLA will be programmed based on one set of functionality or the other. In accordance with an embodiment of the invention, the size of the PLA will be generated according to the larger set of functionality and then both sets of functionality will be considered in determining what spare p-terms are created and how depopulation is performed.

[0072] In still another embodiment, a user may not be as concerned about the size of the PLA, but may want to maintain as much flexibility as possible as long as the PLA remains fast enough and exhibits certain performance characteristics. In such instances, a depopulated PLA could be modelled as described throughout this

document and then tested to measure the performance. If the performance exceeds the selected minimum performance goal (e.g., speed), then spare p-terms are added to the model. The device is then tested again, and the process repeats until the goal is met, e.g., spare p-terms are added until the device just meets or minimally exceeds its minimum acceptable speed characteristics.

* * * * *

SUB
B67

10 [0073] Therefore, a device has been described that allows the formation of smaller, faster, yet flexible PLAs. Unlike conventional “folding” techniques, a device in accordance with the invention requires no regular array structure. Further, distinct from folding, a device in accordance with the invention can depopulate the PLA to a maximum, but then selectively repopulate to allow for future flexibility in reprogramming.

15 [0074] It should be understood that the particular embodiments described above are only illustrative of the principles of the present invention, and various modifications could be made by those skilled in the art without departing from the scope and spirit of the invention. Thus, the scope of the present invention is limited only by the claims that follow.